

A Domain Specific Language for Performance Portable Molecular Dynamics Algorithms

William Robert Saunders^a, James Grant^b, and Eike Hermann Müller^{a,*}

University of Bath, Bath BA2 7AY, Bath, United Kingdom

^a*Department of Mathematical Sciences*

^b*Department of Chemistry*

^{*}*Email: e.mueller@bath.ac.uk*

April 12, 2017

Abstract

Developers of Molecular Dynamics (MD) codes face significant challenges when adapting existing simulation packages to new hardware. In a continuously diversifying hardware landscape it becomes increasingly difficult for scientists to be experts both in their own domain (physics/chemistry/biology) and specialists in the low level parallelisation and optimisation of their codes. To address this challenge, we describe a “Separation of Concerns” approach for the development of parallel and optimised MD codes: the science specialist writes code at a high abstraction level in a domain specific language (DSL), which is then translated into efficient computer code by a scientific programmer. In a related context, an abstraction for the solution of partial differential equations with grid based methods has recently been implemented in the (Py)OP2 library. Inspired by this approach, we develop a Python code generation system for molecular dynamics simulations on different parallel architectures, including massively parallel distributed memory systems and GPUs. We demonstrate the efficiency of the auto-generated code by studying its performance and scalability on different hardware and compare it to other state-of-the-art simulation packages. With growing data volumes the extraction of physically meaningful information from the simulation becomes increasingly challenging and requires equally efficient implementations. A particular advantage of our approach is the easy expression of such analysis algorithms. We consider two popular methods for deducing the crystalline structure of a material from the local environment of each atom, show how they can be expressed in our abstraction and implement them in the code generation framework.

keywords: Molecular Dynamics, Domain Specific Language, Performance Portability, Parallel Computing, GPU

1 Introduction

Molecular Dynamics (MD) codes such as NAMD [1, 2], LAMMPS [3], GROMACS [4, 5] and DL-POLY [6, 7] are important computational tools for understanding the fundamental properties of physical, chemical and biological systems. They can be used to verify phenomenological theories about atomistic interactions, understand complex biomolecules [8] and self assembly processes [9], replace costly laboratory experiments and allow access to areas of parameter space which are very difficult to reproduce experimentally. For example, simulations can be run at high pressures and temperatures found in stellar atmospheres [10], or for dangerous substances, such as radioactive materials (see e.g. [11]). Classical MD codes simulate a material by following the time evolution of a large number of particles which obey the laws of classical physics (in particular Newton’s laws [12]) and interact via phenomenological potentials. To extract meaningful information, the state of the system (i.e. the distribution of particle positions and velocities) has to be analysed, for example by calculating pairwise distribution functions. Information on the crystalline structure of a material can

be derived by inspecting the local environment of each particle [13, 14, 15].

In order to study systems at physically relevant length- and timescales and to produce statistically converged results, modern codes typically run in parallel on state-of-the-art supercomputers [2]. With the recent rise of novel manycore chips, such as GPU and Xeon Phi processors, several popular MD simulation packages have been successfully adapted to those new architectures, see e.g. [16, 17, 18, 19, 20, 21]. However, developers of MD codes face significant challenges: adapting and optimising existing codes requires not only a deep understanding of the physics and chemistry of the simulated system, but also detailed knowledge of the rapidly evolving hardware. To name just a few complications, GPUs have a complex memory hierarchy (host/device memory, shared memory and local registers) and any data access has to be coalesced to avoid unnecessary data movement. Write conflicts have to be avoided in threaded implementations on manycore chips and recent CPUs, such as the Intel Haswell and Broadwell chip, only run at peak performance if the code can be vectorised. Since in practice it is rare for a chemist/physicist to possess the skills for op-

timising code on this level, it can be very challenging to port MD software to a new architecture and maintain its performance in a rapidly evolving hardware landscape. To address this fundamental issue, we describe an approach based on the idea of a “Separation of Concerns” between the domain specialist and scientific programmer. By using a suitable abstraction, both the scientific capabilities and computational performance can be improved independently.

DSLs for grid-based PDE solvers Very similar issues have been faced by developers of grid-based solvers for partial differential equations (PDEs). The key observation there was that the fundamental and computationally most expensive operations can be expressed in terms of a suitable abstraction: the algorithms (e.g. explicit time stepping methods or iterative solvers for elliptic PDEs) can be formulated as the repeated iterations over a set of grid entities (cells, vertices, faces, edges), each of which can hold information, such as a local field value. This expression of the algorithm in a Domain Specific Language (DSL) simplifies the implementation significantly: once the domain-specialist has expressed the code in terms of those basic operations at the correct abstraction level and encapsulated any data in the corresponding fundamental data structures, a computational scientist can implement and optimise the code on a particular architecture.

By introducing the correct abstraction, only a small set of typical loops, which can be parametrised over the set of input and output data, has to be considered. This concept has been applied very successfully in the development of the performance-portable OP2 library [22, 23], which allows the execution of finite element and finite volume codes on a range of architectures. As demonstrated in [24, 22, 23, 25], the code achieves excellent performance on CPUs, GPUs and Xeon Phi processors. Similar techniques for structured grids have been used to develop the C++ based STELLA grid library for the COSMO numerical weather forecast model [26].

Recently OP2 was re-implemented in Python as the PyOP2 [27] framework. In PyOP2 the science user specifies the computationally most expensive operations as a set of small kernels written in C. Using code generation techniques, those kernels are then compiled and executed on a particular architecture. By employing just-in-time compilation, the kernels are launched from a high-level Python code which implements the overall solver algorithm. The performance of the resulting code is on a par with that of monolithic Fortran- or C- implementations.

A new DSL for MD simulations In this paper we describe a similar DSL approach for molecular dynamics simulations. The fundamental operation we consider is a two-particle kernel: the user implements a short C-code which is executed for each combination of particle pairs in the simulation. This kernel can modify any properties stored on those particles. A classic example is the force calculation: for each pair of particles, the force (output) is calculated as a function of the two particle positions




Figure 1: Structure of the code generation framework. The “Separation of concerns” between the domain specialist user and computational scientist is indicated by the dashed horizontal line.

(input). This local operation can be expressed in a few lines of C-code. The code is then executed over all particle pairs, using the optimal algorithm for a particular hardware and the nature and size of the problem. For example, on a CPU architecture, cell-list or neighbour-list methods can be used, whereas on GPU a neighbour-matrix approach as in [28] might be more suitable. Those details of the kernel execution, however, are of no interest for the science developer who can focus on (i) the implementation of the local kernel and (ii) the overall algorithm which orchestrates the kernel calls in an outer timestepping loop.

To achieve this we developed a Python-based code generation system which creates and compiles fast, architecture dependent wrapper code to execute the C-kernel over all particle pairs. Our approach is shown schematically in Fig. 1. By using Python as a high-level language, looping algorithms such as the Velocity Verlet method [29] (see also e.g. [30, 31]) for timestepping or advanced thermostats [32, 33] can be implemented very easily, while still generating fast code for the computationally expensive particle loops.

In the following we describe a proof-of-concept implementation of the abstraction and concentrate on short-range two-particle kernels, i.e. kernels which are only executed for particles which are separated by no more than a specified cutoff distance. We demonstrate that for a Lennard-Jones benchmark we achieve performance similar to state-of-the-art simulation tools such as DL-POLY and LAMMPS. As discussed in Section 6, further generalisations of the framework to support long range electrostatic interactions, multiple species and bonded interactions for molecules will be implemented in the future.

We stress, however, that our approach is not limited

to force calculations. To extract meaningful information from a simulation, the results have to be analysed. With growing problem sizes and data volumes, this step becomes computationally expensive and requires efficient and parallel implementations. Below we consider two methods for analysing local environment which can be used to classify the crystalline phase of a material: the bond order analysis in [13] and common neighbour analysis in [14] (see also [15] for an overview of other analysis methods). In the traditional approach, the user would run the simulation with an existing MD package and then write post-processing code to extract physically meaningful information from the output. However, in contrast to the MD code itself, parallelising this analysis code or porting it to a different architecture is often too time consuming to be feasible. As we will demonstrate below, the fundamental kernels for various common analysis methods can be expressed in our framework. This implies that optimised and parallel code is automatically generated for this important stage of the simulation workflow.

A high-level approach for introducing new algorithms to existing MD packages has been realised in the PLUMED [34] and MIST [35] libraries. They are written as plug-ins to well-established codes and introduce free energy methods and alternative integrators respectively. However, this approach still requires the underlying MD code to be implemented efficiently in the first instance. Similar high-level Python interfaces are provided by OpenMM [36] and HOOMD-blue [17]; in those two cases the underlying code is part of the package itself. Using these interfaces both OpenMM and HOOMD-blue allow the user to control a simulation and access available particle data through calls to the underlying library.

Many MD libraries support the implementation of custom interactions by either providing a mechanism that interpolates tabulated values to produce a potential, or a plugin system that allows users to write and compile extensions that implement the desired interaction. The OpenMM Python interface allows a custom potential to be described in symbolic form. Based in this, OpenMM will automatically generate GPU code by using symbolic differentiation and code generation. The resulting code is compiled at runtime through the OpenCL compiler.

However in all cases (with the exception of kernel code generation in OpenMM) the primary aim of the provided Python interface is to provide access functionality in an underlying C++ or Fortran code, i.e. Python acts as a “glue” for combining existing functionality. If a desired simulation or technique cannot be described within the Python interface for the library, the user needs to program extensions for the specific library. In contrast, our approach is more invasive and allows the expression of both the high-level algorithm and low level kernel in one code. We support general kernels, which are not restricted to force calculations that can be expressed in mathematical form.

Structure This paper is organised as follows: in Section 2 we introduce the fundamental abstractions and data structures used in our approach. The implementa-

tion of the resulting DSL in a Python library and code generation techniques for different architectures are discussed in Section 3. In Section 4 we show how fairly complex structure analysis techniques based on bond order and common neighbour analysis can be expressed in our abstraction and explain how they can be added to the simulation. To demonstrate the performance of the generated code, we compare runtime and scalability to other popular MD packages both on MPI-parallel clusters and for GPUs in Section 5. Here we also show output of the structure analysis algorithms described in Section 4. We conclude and outline ideas for further developments in Section 6.

2 Abstraction

We begin by formulating the key operations which are required to develop a generic MD code. If the domain specialist (computational physicist or chemist) can express their algorithms in terms of those operations, then the code can be implemented in a performance portable way in the high-level Python framework described in Section 3.

Throughout this paper we assume that we want to simulate and analyse a collection of $N \gg 1$ particles. Let each particle with global index $i \in \{0, 1, 2, \dots, N-1\} \equiv \mathcal{N}$ have a set of properties π such that $\pi_r^{(i)}$ is the value of the r -th property on particle i . Each particle has exactly M properties, i.e. $r \in [0, M-1] \equiv \mathcal{M}$. Properties can, for example, be the particle’s position and momentum vector, its charge or the particle index. In addition there can be M^g global properties $\pi_{r^g}^g$ with $r^g \in [0, M^g-1] \equiv \mathcal{M}^g$. Typical global properties might be the total kinetic energy or the radial distribution function (represented as a vector R_i which counts the average number of particles in each distance interval $[r_i, r_{i+1})$).

Operations which involve one or more particles are described in the following three definitions:

Definition 1. A *Particle Loop* is an operation which for each particle $i \in \mathcal{N}$ reads properties $\pi_r^{(i)}$ with $r \in \mathcal{M}_R \subset \mathcal{M}$ and writes properties $\pi_s^{(i)}$ with $s \in \mathcal{M}_W \subset \mathcal{M}$. The operation can also read global properties $\pi_{r^g}^g$ with $r^g \in \mathcal{M}_R^g \subset \mathcal{M}^g$ and write $\pi_{s^g}^g$ with $s^g \in \mathcal{M}_W^g \subset \mathcal{M}^g$ such that the final value of these global properties is independent of the order in which it loops over the particles.

Example 1. *Kinetic energy calculation.* To calculate the total kinetic energy, we loop over all particles i and add $\frac{1}{2}m^{(i)} \sum_{k=0}^{d-1} (v_k^{(i)})^2$ to the global variable K . The particle properties considered in this example are the mass $m^{(i)}$ and the three components $v_k^{(i)}$, $k = 0, 1, 2$ of the particle’s velocity vector $\mathbf{v}^{(i)}$.

Definition 2. A *Particle Pair Loop* is an operation which for all particle pairs $(i, j) \in \mathcal{N} \times \mathcal{N}$ reads properties $\pi_r^{(i)}$ and $\pi_r^{(j)}$ with $r \in \mathcal{M}_R \subset \mathcal{M}$ and modifies properties $\pi_s^{(i)}$ with $s \in \mathcal{M}_W \subset \mathcal{M}$ such that the result is independent of the order of execution. The kernel can also read global properties $\pi_{r^g}^g$ with $r^g \in \mathcal{M}_R^g \subset \mathcal{M}^g$ and

write $\pi_{s^g}^g$ with $s^g \in \mathcal{M}_W^g \subset \mathcal{M}^g$ such that the result does not depend on the order in which the loop is executed over all particle pairs.

Example 2. Force Calculation. The most obvious example of a Particle Pair Loop is the force calculation. Here each particle has six relevant properties, namely the three entries of its position vector and the three entries of the force exerted on the particle by all other particles. For each particle pair the total force on the first particle is incremented by the interaction force $\mathbf{f}(\mathbf{r}^{(i)}, \mathbf{r}^{(j)})$ which depends on the relative position of the particles, i.e. the three position properties $r_k^{(i)}$ for $k = 0, 1, 2$ are read and the three force properties $F_k^{(i)}$ are incremented as $F_k^{(i)} \mapsto F_k^{(i)} + f_k(\mathbf{r}^{(i)}, \mathbf{r}^{(j)})$.

Definition 3. A Local Particle Pair Loop is a Particle Pair Loop which is only executed for particles which are separated by no more than a specified cutoff distance r_c .

Example 3. Local environment. Suppose that each atom can be in one of two possible states. For every atom we want to count the number of other atoms in the same state which are up to a distance r_c away. In this case each particle would have five properties, namely the three entries of the position vector, the state of the atom and the number of atoms in the same state in the local environment. For each pair of atoms the Particle Pair Kernel would first check whether they are less than r_c apart by calculating the distance $|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}|$ between the particle positions. If this is the case, and both particles are in the same state, the counter for the number of same-state atoms is increased.

Further examples will be given in Section 4 where we show how the bond order analysis in [13] and a common neighbour analysis [14] can be expressed as Particle- and Particle Pair- Loops. The Particle Pair Loop can easily be generalised to a loop involving $k > 2$ particles for multi-particle forces.

Note that the computational complexity of a Local Particle Pair loop is $\mathcal{O}(N \cdot N_{\text{local}})$ where $N_{\text{local}} = (4/3)\pi r_c^3 \rho$ is the average number of local neighbours. Since, for constant density ρ , the number N_{local} is constant and relatively small, the computational complexity is $\mathcal{O}(N)$ and therefore significantly smaller than the $\mathcal{O}(N^2)$ complexity of a Particle Loop.

Comment on Newton’s third law For most physically relevant interactions the force on the first particle of the pair is equal and opposite to the force acting on the second particle. Hence, instead of looping over all $N(N-1)$ unordered pairs (i, j) , one could also only loop over the $N(N-1)/2$ ordered pairs with $i < j$, calculate the force once and update it on both particles. Naively this should lead to a speedup of a factor of two. However, it introduces write conflicts in a (shared memory) parallel implementation. While those can be avoided by adding suitable atomic statements or using a colouring approach, the more serious issue is that it prevents vectorisation, which is crucial to obtain optimal performance on modern CPUs. When writing back to memory, the compiler

has to assume that there could be aliasing between particle data (from the compiler’s point of view two of the neighbours of each particle could be identical), and will not generate vectorised code.

As a result this factor of two is more than offset by the advantages of vectorisation and we find that the code is faster overall if we loop over all ordered pairs and only write to the first particle. As will be demonstrated in Section 5.1, for short range forces we achieve equal or better performance than other common MD packages. If necessary, it would of course be possible to implement a version of the pair looping mechanism which exploits Newton’s second law in our code generation framework and this could be considered in future extensions.

3 Implementation

The operations identified in the previous section are the computationally most expensive components of an MD simulation. We now describe their efficient parallel implementation in a code generation framework. From the discussion above it should be clear that our framework will have to provide (1) data structures to represent particle properties $\pi_r^{(i)}$ as well as global properties π_r^g and (2) mechanisms for executing Particle- and Particle Pair-Loops. The following choices are inspired by the PyOP2 [27] data structures and execution model. An implementation of the framework described in this section can be found at:

<https://bitbucket.org/wrs20/ppmd>

All results in this paper were obtained with the release available as [37].

3.1 Data structures

Particle properties $\pi_r^{(i)}$ are represented as instances of a `ParticleDat` class. This class is a wrapper around a two-dimensional `numpy` array, where the first index labels the particle i and the second corresponds to the property index r . Similarly we provide storage for global data shared by all particles in a `ScalarArray` class.

For convenience and to support different data types, we do not collect all properties into a single `ParticleDat` (or `ScalarArray`), but rather allow several `ParticleDats` and `ScalarArrays` instances which can be named by the user. For example, consider a simulation with particles which have a three dimensional position and momentum and vectors $\mathbf{r}^{(i)}, \mathbf{v}^{(i)} \in \mathbb{R}^3$ and a species index $S^{(i)} \in \mathbb{N}$. We also store the total kinetic- and potential energies $KE, PE \in \mathbb{R}$. This set of local and global properties would be implemented as shown in Listing 1.

Listing 1: Data structure initialisation

```
x = ParticleDat(ncomp=3, dtype=float)
v = ParticleDat(ncomp=3, dtype=float)
S = ParticleDat(ncomp=1, dtype=int,
               initial_value=0)
KE = ScalarArray(ncomp=1, dtype=float,
                initial_value=0.0)
PE = ScalarArray(ncomp=1, dtype=float,
```



```
initial_value=0.0)
```

The underlying `numpy` array can be accessed as the `ParticleDat.data` property; however the “getitem” and “setitem” methods have been overloaded to automatically mark the `ParticleDat` as “dirty” if the internal data has been modified directly by the user. This is important in parallel implementations based on a domain decomposition approach, where data owned by neighbouring processors is duplicated in a “halo” region. If “dirty” data is used subsequently in a loop, a exchange of halo data will be triggered automatically and ensures that data is consistent between processors. The interface to the stored data is identical for both CPU- and GPU- `ParticleDat` data structures. When accessing data stored in a `ParticleDat` stored on the GPU in device memory, “getitem” and “setitem” calls will automatically trigger data copies between host- and device-memory. The correct architecture is chosen at the beginning of the Python script by setting aliases for the appropriate objects as shown in Listing 2.

Listing 2: Switching between CPU and GPU implementation

```
import ppmd as md
# Set USE_CUDA to True or False
if not USE_CUDA:
    Data = md.data
    State = md.state.State
    ParticleLoop =
        md.loop.ParticleLoop
    PairLoop =
        md.pairloop.PairLoopNeighbourListNS
    PBC =
        md.domain.BoundaryTypePeriodic()
else:
    Data = ppmd.cuda.cuda_data
    State = ppmd.cuda.cuda_state.State
    ParticleLoop =
        ppmd.cuda.cuda_loop.ParticleLoop
    PairLoop =
        ppmd.cuda.cuda_pairloop. \
            PairLoopNeighbourListNS
    PBC =
        ppmd.cuda.cuda_domain. \
            BoundaryTypePeriodic()

PositionDat = Data.PositionDat
ParticleDat = Data.ParticleDat
ScalarArray = Data.ScalarArray
```

3.2 Particle Pair Loops

In addition to data structures, an execution model is required to launch the computational kernel over all particle pairs. For this, the user writes a brief C-kernel which describes how the properties of the two particles involved in the interaction are modified. In addition, the `ParticleDats` which are operated on have to be passed explicitly to the pair looping mechanism. For each `ParticleDat` an access descriptor describes whether the property is read from or written to. The allowed access descriptors are `READ` (property is only read), `WRITE` (property is only written to), `RW` (property is read and written), `INC` (property is incremented) and `INC_ZERO` (identical to

`INC` except the values are set to zero before the kernel is launched). Since the code generation system does not inspect the C-kernel provided by the user, this information allows the looping system to handle read- and write- access to particle properties in a parallel setting. For example, in a distributed memory implementation, before the execution of the loop halo regions have to be updated for all variables which have a `READ` access descriptor. Similarly, if a particle has `WRITE` or `INC` access, in a threaded implementation write conflicts have to be avoided by generating atomic write statements or employing suitable colouring (see for example the layer algorithm described in [38]). In addition to `ParticleDats`, global variables (represented `ScalarArrays`) can be passed to the kernel with the same access descriptors. To treat numerical constants which do not change during the kernel execution, each kernel can also be passed a list of `Constant` objects. Any instances of `Constant` variables in a kernel are replaced by their numerical values at compile time; this allows the compiler to make additional optimisations, for example by exploiting static loop bounds.

As an (fictitious) example, imagine that on each particle we store the properties a (which has $d = 3$ components) and b (which has one component). For all particles i we carry out the operation which calculates

$$b^{(i)} = \sum_{\text{all pairs } (i, j)} \sum_{r=0}^{d-1} (a_r^{(i)} - a_r^{(j)})^2 \quad (1)$$

and updates the global sum

$$S^g = \sum_{\text{all pairs } (i, j)} \sum_{r=0}^{d-1} (a_r^{(i)} - a_r^{(j)})^4. \quad (2)$$

A Particle Pair loop which performs this operation can be implemented as shown in Listing 3. The execution over all particle pairs is illustrated schematically in Fig. 2. Inside the pairloop the two involved particles are accessed as the `.i` and `.j` component of a structure, and the names of the `ParticleDats` are given in the dictionary which is passed as the second argument to the `PairLoop` constructor. For example, the r -th component of the first particle is accessed as `a.i[r]`. This C-variable automatically points to the correct position in the `numpy` array which holds the `ParticleDat` values. Particle Loops are conceptually very similar and can be implemented in the same way.

We note that the code in Listing 3 resembles what would be written in PyOP2 to implement a loop over a set of mesh entities. In PyOP2 the fundamental data types are called `Dat` and `GlobalDat`. A `Dat` object represents data which is associated with topological entities of the mesh, for example the average value of a field in each grid cell. A `GlobalDat` variable contains globally available data. The main difference is that PyOP2 loops over a particular static set of topological entities and can access data on other related entities which are specified via indirection maps. Those indirection maps are provided as additional arguments to the `Dat` dictionary of the looping class. An important difference is that the indirection maps in PyOP2 have a fixed “arity”, i.e. each

Listing 3: Python code for executing the operations in Eqs. (1) and (2) over all particle pairs.

```
# dimension
dimension=3

# number of particles
npart=1000

# Define Particle Dats
a = ParticleDat(npart=npart,
                ncomp=dimension,
                dtype=float)
b = ParticleDat(ncomp=1,
                npart=npart,
                initial_value=0.0,
                dtype=float)
S = ScalarArray(ncomp=1,
                initial_value=0.0,
                dtype=float)

kernel_code='''
double da_sq = 0.0;
for (int r=0;r<dimension;++r) {
    double da = a.i[r]-a.j[r];
    da_sq += da*da;
}
b.i[0] += da_sq;
S += da_sq*da_sq;
'''

# Define constants passed to kernel
kernel_consts = (Constant('dimension',
                           dimension),)

# Define kernel
kernel = Kernel('update_b',
                kernel_code,
                kernel_consts)

# Define and execute pair loop
pair_loop = PairLoop(kernel=kernel,
                      {'a':a(access.READ),
                       'b':b(access.INC),
                       'S':S(access.INC)})

pair_loop.execute()
```

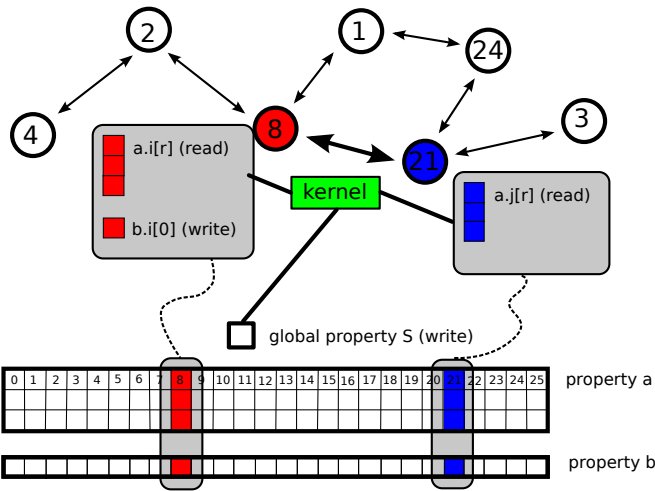


Figure 2: Pairwise kernel for executing the operation in Eqs. (1) and (2) over all particle pairs.

unknown depends on a fixed number of other unknowns. In contrast, in an MD code, the number and identity of nearest neighbours of each particle varies throughout the simulation. In a parallel MD code the distribution of particles over processors also changes over time, and this requires additional parallel communication.

3.3 Code generation for performance-portability

To execute a pairloop we use a code generation approach. Given the kernel and information on how data is accessed, appropriate wrapper C-code for launching the kernel over all particle pairs is generated for a particular hardware backend. This means that to target different architectures, the user has to write the kernel code only once: it is up to the code generation system (developed by a computational scientist) to execute this on a specific architecture. Note also that the user never has to explicitly add calls to MPI routines or guarantee the correctness of the results on a threaded architecture by protecting write statements with “atomic” or “critical” keywords. Similarly, on a particular architecture different pair looping mechanisms (described below) lead to the same scientific result but can have different computational performance. Our method allows the straightforward comparison between different looping mechanism without the user intervention to modify code, a feature that could potentially be exploited to optimise performance on a problem-by-problem basis. Since the system is aware of data dependencies between different kernel, loop fusion to reduce the amount of data movement could be implemented to further improve performance.

On a sequential machine, the simplest possible wrapper code is shown in Listing 4.

Listing 4: Pair loop in a sequential implementation

```
for (int i=0;i<npart;++i) {
    for (int j=0;j<npart;++j) {
        if (i!=j) {
            // INSERT KERNEL CODE HERE
        }
    }
}
```

The computational complexity of this nested loop is $\mathcal{O}(N^2)$ and for short range kernels this method would be extremely inefficient. In the following we describe more advanced looping mechanisms for executing Local Particle Pair Loops on parallel architectures.

3.4 Cell based methods for Local Particle Pair Loops

If we only consider Local Particle Pair kernels with a fixed cutoff r_c , the computational complexity is reduced to $\mathcal{O}(N)$ and it is possible to use cell based looping methods (see [39] for an introduction). In this approach the physical domain of size $[0, L_x] \times [0, L_y] \times [0, L_z]$ is divided into small cells of size $\Lambda_x \times \Lambda_y \times \Lambda_z$ such that $\Lambda_{x,y,z} \geq r_c$; to simplify the presentation, we assume $\Lambda = \Lambda_x = \Lambda_y = \Lambda_z$

Listing 5: Creation of a state object with position, velocity and acceleration data

```
# create a State object
state = State()
state.domain = BaseDomain()
state.domain.boundary_condition =
    BoundaryTypePeriodic()

# add ParticleDats to state
state.npart = N
state.pos = PositionDat(ncomp=3,
                        dtype=float)
state.vel = ParticleDat(ncomp=3,
                        dtype=float)
state.acc = ParticleDat(ncomp=3,
                        dtype=float)
```

in the following. At a given point in time every particle can be uniquely associated with one of those small cells. The local Particle Pair loop with cutoff r_c can then be executed by visiting all cells e in an outer loop and then iterating over all 26 neighbouring cells e' . Since $\Lambda \geq r_c$ it is then sufficient to consider pairs of particles (i, j) such that $i \in e$ and $j \in e'$.

For each particle this algorithm considers potential interactions with other particles in a volume $27\Lambda^3$. However, most of these pairs will be separated by a distance $|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}| > \Lambda \geq r_c$. To avoid unnecessary execution of the kernel for non-interacting particles, it is possible to add another preprocessing step which loops through all potential pairs and only stores those which are a distance of up to Λ away. Interactions can then be calculated by looping through this neighbour list. In three dimensions this reduces the cost of the force calculation by up to a factor $81/(4\pi) \approx 6.45$. The computational overhead for building the neighbour list is usually amortised by the gain in the force calculation.

For both $\mathcal{O}(N)$ pair looping mechanisms described above, different **ParticleDats** can no longer be considered independently, but rather have to be seen as members of a **State** object which also stores the shared cell- and neighbour lists. One particular **ParticleDat** in this **State** object stores the particle position, and this information is required when building the cell- or neighbour-lists. To distinguish it from other properties such as velocity and acceleration, a special derived class **PositionDat** is used. As shown in Listing 5, all **ParticleDats** in a simulation have to be associated with a **State** object by setting (user-defined) properties of the state as

```
A.PROPERTY = ParticleDat(...).
```

Each state also contains a domain object, which stores information about the physical domain size and boundary conditions.

During the simulation, particles will move between cells and hence if $\Lambda = r_c$ the cell- and neighbour lists need to be rebuilt at every iteration, which can be very expensive. This can be avoided by increasing the cell size and choosing an extended cutoff \bar{r}_c : if the relevant

interaction range is $r_c < \bar{r}_c < \Lambda$ and v_{\max} is the maximal particle velocity, a rebuild of the cell- and neighbour lists is only necessary every n time steps if

$$\bar{r}_c = r_c + 2n \cdot \delta t \cdot v_{\max} = r_c + \delta. \quad (3)$$

For time integration loops we further provide the **IntegratorRange** class, which allows timestepping methods to be implemented in a way that retains the simplicity and flexibility of a standard Python **range** based loop without explicit cell- and neighbour list rebuilds by the user. An example is shown in Listing 6. In addition to the number of integration steps, **IntegratorRange** is passed the following information:

- the timestep size δt ,
- a **ParticleDat** containing particle velocities,
- a maximum reuse count and
- the thickness $\delta = \bar{r}_c - r_c$ of the additional shell.

Listing 6: Example use of **IntegratorRange** called with: N_i number of iterations, timestep size δt , velocities \mathbf{v} , list reuse count N_s and shell thickness $\delta = \bar{r}_c - r_c$.

```
for i in IntegratorRange(Ni, dt, v,
                        Ns, delta):
    particle_loop_1.execute()
    force_calculation.execute()
    particle_loop_2.execute()
```

3.4.1 Parallelisation

To simulate the interactions of a very large number of particles in a reasonable time, MD codes have to be parallelised. Modern HPC installations expose parallelism on different levels and the implication of this complex hierarchy on MD implementations will be discussed in the following.

Distributed memory The cell-based methods described above can be parallelised with a standard domain-decomposition approach. For this the global domain is split up into smaller subdomains stored on each processor. To correctly include interactions with particles stored on neighbouring subdomains, a layer of halo cells is added. Those cells hold copies of particles which are *owned* by other processors. Data in halo cells needs to be updated whenever this data changes. Note, however, that this is only necessary if the values of a particular **ParticleDat** are actually read in the loop, and this information is made explicit via the access descriptors passed to the pairloop. Our code generation system will therefore only launch the corresponding parallel communication calls if necessary. This guarantees the parallel correctness of the code while avoiding superfluous and expensive parallel communications. Since particles can move to different processors and hence the data layout is not fixed, there are actually two parallel communication types which need to be carried out:

1. Data on particles in the halo region has to be updated if it has changed and is used in a pair loop.

2. Particles which leave the local domain need to be moved to a different processor.

The first operation typically needs to be performed whenever dirty data is read. For example halo exchanges on particle positions are required before every force update. The second communication type only needs to be performed every n steps, since the increased cut off in Eq. (3) ensures the accuracy of the calculation even if a particle leaves the cell during those steps. In addition to rebuilding the cell list, when a particle has left the sub-domain owned by a processor, the `State` object will automatically move all data owned by the particle to the receiving processor.

Virtually all modern supercomputers now consist of a large collection of relatively complex compute units (CPUs, GPUs or Xeon Phis) organised into nodes. While parallelisation between nodes is achieved with the distributed memory approach described above, each node consists of a large number of compute cores which have access to the same memory. Parallelisation across those cores on a node requires a different approach which will be described in the following section. To make use of the full machine, a hybrid approach which combines both parallelisation strategies is typically used.

Threading and GPU parallelisation To reduce memory requirements, in a sequential implementation (or if the code is parallelised purely with a distributed memory approach), the cell-list is stored as a linked list and the neighbour list is realised by storing all neighbours in a long array. This prevents any further shared memory vectorisation based on threading since neither the cell-list nor the neighbour-list can be built in parallel. To avoid this problem on GPUs we use the approach in [28] and replace the cell list by a cell-occupancy matrix H . For this each particle i is associated with a cell c_i and the particles in a cell are arranged into “layers”, such that all particles in a cell have a different layer-index. If the layer index of particle i is ℓ_i , then $H_{c_i, \ell_i} = i$, and H can be built in parallel. Based on this, a neighbour matrix W can be built such that $W_{m,i}$ is the index of the m -th neighbour of particle i . An alternative approach which is described in [38] and avoids building W , would be to loop over all pairs of layers and use the matrix H to identify interacting particles.

Vectorisation Modern HPC CPUs contain floating point units (FPU) which are capable of executing Single Instruction Multiple Data (SIMD) instructions. By using SIMD instructions the FPU can apply the same operation to multiple data points simultaneously. For example a 256bit wide vector FPU may simultaneously apply the same operation to four 64bit doubles or eight 32bit floats. However, producing machine code that contains these SIMD instructions is a non-trivial task. One approach to produce SIMD instructions involves explicitly implementing the desired mathematical operations using “intrinsic” functions for a target architecture. This ensures that SIMD instructions are generated by the compiler but requires careful implementation to be technically correct

and produce efficient code. Since the intrinsics are hardware specific, this approach is not portable. In our code we currently simply avoid code patterns which inhibit auto-vectorisation by the compiler. The ability to replace loop bounds by their numerical values via `Constant` objects also helps with vectorisation. As noted in Section 2, we do not currently exploit symmetry in Newton’s third law when computing forces between particles (although the framework would in principle support this). We find that the Intel C/C++ compiler will successfully auto-vectorise kernels without explicitly implementing gather or scatter operations providing the kernel itself does not contain a code pattern that inhibits vectorisation. The strong- and weak- scaling results reported in Section 5.1 were obtained with vectorised code. In the future we will also explore further optimisations which are necessary for more complex kernels.

4 Structure analysis algorithms

To demonstrate that the abstraction and implementation described in the previous sections can be used to implement more complex kernels and is not restricted to force calculations, we now discuss two popular algorithms for classifying the local environment of a particle. We show how these algorithms can be expressed in terms of particle- and local particle-pair loops. Both algorithms can be used to identify the crystalline structure of the material; an overview of other common methods can be found in [15].

4.1 Bond order analysis

The bond order analysis (BOA) in [13] introduces a set of order parameters which are defined for each particle i as

$$Q_\ell^{(i)} = \sqrt{\frac{4\pi}{2\ell+1} \sum_{m=-\ell}^{+\ell} |q_{\ell m}^{(i)}|^2}. \quad (4)$$

The sum

$$q_{\ell m}^{(i)} = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} Y_\ell^m(\hat{\mathbf{r}}^{(i,j)}) \quad (5)$$

is computed by evaluating the spherical harmonics Y_ℓ^m in the directions

$$\hat{\mathbf{r}}^{(i,j)} = \frac{\mathbf{r}^{(i)} - \mathbf{r}^{(j)}}{|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}|}$$

pointing from the atom i to each of its neighbours $j \in \mathcal{N}(i)$. Atoms are considered to be neighbours if their distance is smaller than a predefined cutoff range r_c . The moments $q_{\ell m}^{(i)}$ describe the angular dependence of the charge density $\rho^{(i)}(\mathbf{r} - \mathbf{r}^{(i)})$ of the atom’s neighbours in spectral space. It can then be shown that the integral of the squared averaged charge density can be written as

$$\int_\Omega |\rho^{(i)}(\mathbf{r})|^2 d\Omega = \sum_{\ell=0}^{\infty} \left(Q_\ell^{(i)} \right)^2.$$

lattice structure	Q_4	Q_5	Q_6
fcc	0.191	0	0.575
hcp	0.097	0.252	0.485
bcc	0.036	0	0.511

Table 1: Values of Q_4 and Q_6 for perfect lattices, see [15] and Tab. 1 in [40].

Perfect crystal lattices have well defined values for $Q_\ell = Q_\ell^{(i)}$ which are the same for all atoms. In particular the order parameters with $\ell = 4, 5, 6$ are often used to estimate the degree and nature of crystallinity. Specific values for fcc, hcp and bcc lattices are given in Tab. 1 ([15, 40]). In a simulation the local structure of the material can therefore be estimated by calculating $Q_\ell^{(i)}$ and comparing to the reference values in Tab. 1. If they agree within some tolerance, the system is classified to be in the corresponding state.

The order parameters $Q_\ell^{(i)}$ can be calculated with the two loops shown in Algorithms 1 and 2. The first particle pair loop (Algorithm 1) calculates the number of neighbours $\nu_{\text{nb}}^{(i)} = |\mathcal{N}(i)|$ and the moments

$$\tilde{q}_{\ell m}^{(i)} = \sum_{j \in \mathcal{N}(i)} Y_\ell^m(\hat{\mathbf{r}}^{(i,j)}) \quad (= \nu_{\text{nb}}^{(i)} q_{\ell m}^{(i)})$$

for $m = -\ell, \dots, +\ell$ for each atom i ; those quantities are stored in two `ParticleData`s. The particle loop in Algorithm 2 uses $\nu_{\text{nb}}^{(i)}$ and $\tilde{q}_{\ell m}^{(i)}$ to calculate the $Q_\ell^{(i)}$ according to Eq. (4); the result is stored in a third `ParticleData`. The corresponding source code can be found in the `examples/structure/boa/` subdirectory of the accompanying code release [37].

Algorithm 1 BOA Local Particle Pair Loop I.

Input: particle positions $\mathbf{r}^{(i)}$ [READ].

Output: moments $q_{\ell m}^{(i)}$ [INC_ZERO]

```

1: for all pairs  $(i, j)$  do
2:   if  $|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}| < r_c$  then
3:      $\hat{\mathbf{r}}^{(i,j)} \mapsto (\mathbf{r}^{(i)} - \mathbf{r}^{(j)})/|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}|$ 
4:     for  $m = -\ell, \dots, +\ell$  do
5:        $\tilde{q}_{\ell m}^{(i)} \mapsto \tilde{q}_{\ell m}^{(i)} + Y_\ell^m(\hat{\mathbf{r}}^{(i,j)})$ 
6:     end for
7:   end if
8: end for
```

4.2 Common neighbour analysis

Common neighbour analysis (CNA) [14] is a purely topological method for classifying the local environment of each particle. All atoms within a certain cutoff distance r_c are considered to be “bonded”. For any bonded pair (i, j) the set of all other atoms which are bonded to both i and j are referred to as *common neighbours*. The bonds between those common neighbours define a graph \mathcal{G} . For each pair $(i, j) \in \mathcal{G}$ this graph is now classified by three

Algorithm 2 BOA Particle Loop II.

Input: moments $\tilde{q}_{\ell m}^{(i)}$ [READ], number of local neighbours $\nu_{\text{nb}}^{(i)}$ [READ].

Output: $Q_\ell^{(i)}$ [WRITE]

```

1: for all particles  $i$  do
2:   for  $m = -\ell, \dots, +\ell$  do
3:      $q_{\ell m}^{(i)} \mapsto \tilde{q}_{\ell m}^{(i)} / \nu_{\text{nb}}^{(i)}$ 
4:   end for
5:    $Q_\ell^{(i)} \mapsto \sqrt{\frac{4\pi}{2\ell+1} \sum_{m=-\ell}^{+\ell} |q_{\ell m}^{(i)}|^2}$ 
6: end for
```




Figure 3: Common neighbour analysis for bonded atom pair (i, j) (empty circles). The set of common neighbours (filled circles) are classified as a $(4, 2, 1)$ triplet.

numbers [15]: (1) the number of common neighbours n_{nb} , i.e. the number of vertices in \mathcal{G} , (2) the number of bonds n_{b} , i.e. the number of edges in \mathcal{G} , and (3) n_{lcb} , the number of bonds in the largest cluster (connected subgraph) $\mathcal{G}' \subset \mathcal{G}$. For each pair of bonded atoms this defines a triplet $(n_{\text{nb}}, n_{\text{b}}, n_{\text{lcb}})$ (see Fig. 3). To classify the local environment of an atom, the triplets $(n_{\text{nb}}, n_{\text{b}}, n_{\text{lcb}})$ are computed for all its neighbours and compared to reference signatures for periodic crystal structures. For example, in an hcp lattice, each atom has 12 bonds, six of which are classified as $(4, 2, 1)$ and the other six are $(4, 2, 2)$; see Tab. 1 in [15]. There is some ambiguity in the cutoff distance r_c . To overcome this limitation, the author of [15] suggests an adaptive extension of the method. While this improved algorithm can also be implemented in our framework, for the sake of brevity we do not discuss this extension here and focus on the original method.

To implement the CNA algorithm in our framework we proceed in two steps: For each atom i we first calculate all directly and indirectly bonded atoms. The set $\mathcal{E}_d^{(i)}$ describes the direct bonds; the indirect bonds in the local environment are collected in $\bar{\mathcal{E}}^{(i)}$ (see Fig. 4):

$$\begin{aligned} \mathcal{E}_d^{(i)} &= \{(i, v) : v \in \mathcal{N}, |\mathbf{r}^{(i)} - \mathbf{r}^{(v)}| < r_c\} \\ \bar{\mathcal{E}}^{(i)} &= \{(v, w) : v, w \in \mathcal{N}, |\mathbf{r}^{(v)} - \mathbf{r}^{(w)}| < r_c, \\ &\quad |\mathbf{r}^{(i)} - \mathbf{r}^{(v)}| < r_c\} \end{aligned} \quad (6)$$

Since some of the indirect bonds are counted twice in $\bar{\mathcal{E}}^{(i)}$, the set $\mathcal{E}^{(i)}$ is an ordered representation of the same




Figure 4: Example of direct (left) and indirect (centre and right) bonds as described by the sets $\mathcal{E}_d^{(i)}$, $\bar{\mathcal{E}}^{(i)}$ and $\mathcal{E}^{(i)}$ in Eqns. (6) and (7). The bond (v, w) in the central diagram would be counted twice in $\bar{\mathcal{E}}^{(i)}$ but only once in $\mathcal{E}^{(i)}$.

bonds:

$$\mathcal{E}^{(i)} = \{(v, w) : (v, w) \in \bar{\mathcal{E}}^{(i)}, v < w\} \subset \bar{\mathcal{E}}^{(i)} \quad (7)$$

As before, $\mathcal{N} = \{0, \dots, N-1\}$ is global index set and $\mathcal{N}(i)$ the set of all neighbours of particle i , i.e. all other particles which are no more than a distance r_c away. In a second step we loop over all pairs (i, j) of atoms and calculate the sets

$$\begin{aligned} \mathcal{C} &= \mathcal{N}(i) \cap \mathcal{N}(j) \\ \mathcal{E} &= \{(v, w) : v, w \in \mathcal{C}, v < w\} \subset \mathcal{E}^{(i)} \cap \mathcal{E}^{(j)}. \end{aligned} \quad (8)$$

\mathcal{C} is the set of common neighbours and \mathcal{E} is the set of common neighbour bonds. Note that, to avoid double counting, here we consider ordered bounds $(v, w) \in \mathcal{E}^{(i)}$ such that $v < w$. Together the two sets \mathcal{C} and \mathcal{E} define the graph \mathcal{G} introduced above. The first two entries of the triplet (n_{nb}, n_b, n_{lcb}) can be calculated directly as $n_{nb} = |\mathcal{C}|$ and $n_b = |\mathcal{E}|$. To calculate the size of all subgraphs $\mathcal{G}' \subset \mathcal{G}$, a random node $v \in \mathcal{G}$ is chosen. The size of the subgraph \mathcal{G}' such that $v \in \mathcal{G}'$ is obtained with a breadth-first traversal of the connected component containing v , removing all visited nodes from \mathcal{G} in the process. This is repeated until all nodes have been removed, thus calculating the size of all subgraphs $\mathcal{G}' \subset \mathcal{G}$. The computation of the maximal cluster size $n_{lcb} = \max_{\mathcal{G}' \subset \mathcal{G}} \{|\mathcal{G}'|\}$ with this method is shown explicitly in algorithm 7 in C.

We now show how the CNA algorithm can be implemented as a set of Local Particle Pair- loops. For this, define the following **ParticleDats**:

- **\mathbf{r} (ncomp=3)**: Particle coordinates, $\mathbf{r}^{(i)}$ stores the position of particle i
- **\mathbf{G} (ncomp=1)**: Global id, $G^{(i)} = i \in \mathcal{N}$ stores the unique global index of particle i .
- **ν_{nb} (ncomp=1)**: Number of neighbours, i.e. $\nu_{nb}^{(i)} = |\mathcal{N}(i)|$; this is the number of red particles in the inner circle in Fig. 5.
- **ν_b (ncomp=1)**: Number of bonds in the local environment. $\nu_b^{(i)} = |\mathcal{E}_d^{(i)} \cup \bar{\mathcal{E}}^{(i)}|$ counts the directly bonded neighbours of a particle plus the number of indirect bonds defined in Eq. (6).
- **\mathbf{E} (ncomp= $2\nu_b^{(max)}$)**: Array representation of the set $\mathcal{E}_d^{(i)} \cup \bar{\mathcal{E}}^{(i)}$ defined in Eq. (6). Two consecutive entries $E_{2k}^{(i)}, E_{2k+1}^{(i)}$ represent a bonded pair in the local




Figure 5: Local bonds used for CNA construction

environment of particle i , i.e. one of the links shown in Fig. 5. The entries of $E^{(i)}$ are arranged as follows:

- $(E_{2k}^{(i)}, E_{2k+1}^{(i)}) = (G^{(i)}, G^{(j)})$ with $j \neq i$ for $0 \leq k < \nu_{nb}^{(i)}$
- $(E_{2k}^{(i)}, E_{2k+1}^{(i)}) = (G^{(j')}, G^{(j'')})$ with $j' \neq i, j'' \neq i$ for $\nu_{nb}^{(i)} \leq k < \nu_b^{(i)}$

In other words, the first $\nu_{nb}^{(i)}$ tuples represent the bonds in $\mathcal{E}_d^{(i)}$ and are shown as red (solid) lines in Fig. 5. The remaining $\nu_b - \nu_{nb}$ tuples describe the set $\bar{\mathcal{E}}^{(i)}$ and correspond to the blue (dashed) lines. The static size $\nu_b^{(max)}$ of the list has to be chosen sufficiently large, i.e. $\nu_b^{(max)} \geq \max_i \{\nu_b^{(i)}\}$.

- **\mathbf{T} (ncomp= $3\nu_{nb}^{(max)}$)** stores the triplets (n_{nb}, n_b, n_{lcb}) such that $(T_{3j}^{(i)}, T_{3j+1}^{(i)}, T_{3j+2}^{(i)})$ is the triplet (n_{nb}, n_b, n_{lcb}) for the j -th bonded neighbour of particle i . The number of components $\nu_{nb}^{(max)}$ has to be chosen such that $\nu_{nb}^{(max)} \geq \max_i \{\nu_{nb}^{(i)}\}$.
- **\mathbf{t} (ncomp=1)** stores the number of classified bonds of particle i .

Using those **ParticleDats**, for each particle the list representation $E^{(i)}$ of the set $\mathcal{E}_d^{(i)} \cup \bar{\mathcal{E}}^{(i)}$ can now be calculated with two Local Particle Pair Loops: the first loop, shown in Algorithm 3, calculates the first $2\nu_{nb}^{(i)}$ entries of $E^{(i)}$ by inspecting the direct neighbours of each particle. Based on this, the second loop in algorithm 4 adds the remaining $2(\nu_b^{(i)} - \nu_{nb}^{(i)})$ entries, i.e. the blue (dashed) lines in Fig. 5. The final Particle Pair Loop in algorithm 5 then uses the information stored in $E^{(i)}$ and $E^{(j)}$ to extract the tuple (n_{nb}, n_b, n_{lcb}) .

The C-code for Algorithms 3 and 4 is shown in A.2. All source code (include the one for the slightly longer Algorithm 5) can be found in the subdirectory **examples/structure/cna** of the accompanying code release [37]. Results obtained with our implementation of both a bond order- and common-neighbour-analysis algorithm are shown below in Section 5.2.

5 Results

To demonstrate the performance, portability and scalability of our code generation framework on two different

Algorithm 3 CNA Local Particle Pair Loop I: Calculate direct bonds for each particle.

Input: $\mathbf{r}^{(i)}$ [READ], $G^{(i)}$ [READ].

Output: $\nu_{\text{nb}}^{(i)}$ [INC_ZERO], $\nu_{\text{b}}^{(i)}$ [INC_ZERO], $E^{(i)}$ [WRITE]

```

1: for all pairs  $(i, j)$  do
2:   if  $|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}| < r_c$  then
3:      $(E_{2\nu_{\text{b}}^{(i)}}, E_{2\nu_{\text{b}}^{(i)}+1}) = (G^{(i)}, G^{(j)})$ 
4:      $\nu_{\text{b}}^{(i)} \mapsto \nu_{\text{b}}^{(i)} + 1$ 
5:      $\nu_{\text{nb}}^{(i)} \mapsto \nu_{\text{nb}}^{(i)} + 1$ 
6:   end if
7: end for

```

Algorithm 4 CNA Local Particle Pair Loop II: Calculate all other bonds in the local environment.

Input: $\mathbf{r}^{(i)}$ [READ], $G^{(i)}$ [READ], $\nu_{\text{nb}}^{(i)}$ [READ].

Output: $\nu_{\text{b}}^{(i)}$ [INC], $E^{(i)}$ [RW]

```

1: for all pairs  $(i, j)$  do
2:   if  $|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}| < r_c$  then
3:     for  $k = 0, \dots, \nu_{\text{nb}}^{(j)} - 1$  do
4:       if  $E_{2k+1}^{(j)} \neq G^{(i)}$  then
5:          $(E_{2\nu_{\text{b}}^{(i)}}, E_{2\nu_{\text{b}}^{(i)}+1}) = (E_{2k}^{(j)}, E_{2k+1}^{(j)})$ 
6:          $\nu_{\text{b}}^{(i)} \mapsto \nu_{\text{b}}^{(i)} + 1$ 
7:       end if
8:     end for
9:   end if
10: end for

```

Algorithm 5 CNA Local Particle Pair Loop III: Calculate number of common neighbours $n_{\text{nb}}^{(i)}$, number of bonds $n_{\text{b}}^{(i)}$ between those common neighbours and the largest clustersize $n_{\text{lcb}}^{(i)}$.

Input: $\mathbf{r}^{(i)}$ [READ], $\nu_{\text{nb}}^{(i)}$ [READ], $\nu_{\text{b}}^{(i)}$ [READ], $E^{(i)}$ [READ].

Output: $T^{(i)}$ [WRITE], $t^{(i)}$ [INC_ZERO]

```

1: for all pairs  $(i, j)$  do
2:   if  $|\mathbf{r}^{(i)} - \mathbf{r}^{(j)}| < r_c$  then
3:     Set  $\mathcal{C}$  of common neighbours:
4:      $\mathcal{C} \mapsto \{v : \exists k < \nu_{\text{nb}}^{(i)}, \ell < \nu_{\text{nb}}^{(j)}, v = E_{2k+1}^{(i)} = E_{2\ell+1}^{(j)}\}$ 
5:     Construct set  $\mathcal{E}$  of common neighbour bonds:
6:      $\mathcal{E} \mapsto \{\}$ 
7:     for  $k = \nu_{\text{nb}}^{(i)}, \dots, \nu_{\text{b}}^{(i)} - 1$  do
8:       if  $E_{2k}^{(i)} \in \mathcal{C}$  and  $E_{2k+1}^{(i)} \in \mathcal{C}$  then
9:          $(v, w) = (E_{2k}^{(i)}, E_{2k+1}^{(i)})$ 
10:        if  $w > v$  then
11:          swap  $v \leftrightarrow w$ 
12:        end if
13:        if  $(v, w) \notin \mathcal{E}$  then
14:           $\mathcal{E} \mapsto \mathcal{E} \cup (v, w)$ 
15:        end if
16:      end if
17:    end for
18:     $T_{3t^{(i)}+2}^{(i)} \mapsto |\mathcal{C}|$ 
19:     $T_{3t^{(i)}+1}^{(i)} \mapsto |\mathcal{E}|$ 
20:    Calculate largest cluster size, see Algorithm 7:
21:     $T_{3t^{(i)}+2}^{(i)} \mapsto \text{maxClustersize}(\mathcal{E})$ 
22:     $t^{(i)} \mapsto t^{(i)} + 1$ 
23:  end if
24: end for

```

chip architectures, we implemented the Velocity Verlet integrator [29] (see also e.g. [30, 31]) shown in Algorithm 6. Access descriptors for all loops are given in Tab. 2. The main time stepping loop is realised with an `IntegratorRange` iterator (see Section 3.4), which takes care of cell-list and neighbour-list updates. C-kernels for the particle-loops that update velocity and position in lines 6 and 8 are shown in Listings 7 and 8.

Listing 7: velocity and position update kernel in the Velocity Verlet Algorithm 6 (line 6). The constants `dt` and `dht_iMass` are set to δt and $\delta t/(2m)$ and passed to the pairloop as `Constant` objects.

```
v.i[0] += F.i[0]*dht_iMASS;
v.i[1] += F.i[1]*dht_iMASS;
v.i[2] += F.i[2]*dht_iMASS;
r.i[0] += dt*v.i[0];
r.i[1] += dt*v.i[1];
r.i[2] += dt*v.i[2];
```

Listing 8: velocity update kernel in the Velocity Verlet Algorithm 6 (line 8). As in Listing 7, the quantity $\delta t/(2m)$ is passed to the pairloop as a `Constant` object.

```
v.i[0] += F.i[0]*dht_iMASS;
v.i[1] += F.i[1]*dht_iMASS;
v.i[2] += F.i[2]*dht_iMASS;
```

We simulated a Lennard-Jones liquid system of non-bonded particles interacting via the potential

$$V(r) = 4\epsilon \left(\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 + \frac{1}{4} \right) \quad (9)$$

with a specified cutoff r_c . The C-kernel for the calculation of the resulting short-range force in line 7 is given in A.1. The full source code can be found in the `code/examples/lennard-jones` subdirectory of [37]. It should be stressed that exactly the same code can be used to run the simulation both on a CPU and a GPU if the appropriate definitions shown in listing 2 are added at the beginning of the Python code.

Algorithm 6 Velocity Verlet integrator used in Section 5. The system is integrated numerically with a time step of size δt until the final time $T = N_{\max}\delta t$.

- 1: Create `ParticleDats` for forces \mathbf{F} and velocities \mathbf{v} .
- 2: Create `PositionDat` for particle positions.
- 3: Initialise particle positions and velocities.
- 4: Collect `ParticleDats` and `PositionDat` in a `State` object
- 5: **for** timestep $i = 1, \dots, N_{\max}$ **do**
- 6: For all particles i : $\mathbf{v}^{(i)} \mapsto \mathbf{v}^{(i)} + \frac{\delta t}{2m} \mathbf{F}^{(i)}$, $\mathbf{r}^{(i)} \mapsto \mathbf{r}^{(i)} + \delta t \mathbf{v}^{(i)}$
- 7: For all pairs (i, j) : $\mathbf{F}^{(i)} \mapsto \mathbf{F}^{(i)} + \mathbf{f}(\mathbf{r}^{(i)}, \mathbf{r}^{(j)})$
- 8: For all particles i : $\mathbf{v}^{(i)} \mapsto \mathbf{v}^{(i)} + \frac{\delta t}{2m} \mathbf{F}^{(i)}$
- 9: **end for**

line	loop type	access descriptor
6	<code>ParticleLoop</code>	\mathbf{v} [INC], \mathbf{r} [INC], \mathbf{F} [READ], m [READ]
7	<code>ParticlePairLoop</code>	\mathbf{F} [INC.ZERO], \mathbf{r} [READ]
8	<code>ParticleLoop</code>	\mathbf{v} [INC], \mathbf{F} [READ], m [READ]

Table 2: Access descriptors for the loops in the Velocity Verlet Algorithm 6.

parameter	value
Number of atoms: N	10^6
Number of time steps: N_{\max}	10^4
Non-dimensionalised density: ρ	0.8442
Force cutoff: r_c	2.5
Force extended cutoff: $\bar{r}_c = r_c + \delta$	2.75
Steps between neighbour list update:	20^\dagger

Table 3: Parameters of Lennard-Jones benchmark for the strong scaling experiment; units are chosen such that $\sigma = \epsilon = 1$ († = excluding DL-POLY, see main text).

5.1 Comparison to other codes

To verify that the code generation approach does not introduce any sizable computational overheads, we compare the performance of our code to monolithic C/Fortran implementations in well established and optimised MD libraries. For this we performed the same strong scaling experiment with DL-POLY (version 4.08), LAMMPS (release dated 1st March 2016) and our code generation framework (subdirectory `release` of [37]). Raw results can be found in the accompanying data repository [41].

All codes were built with the Intel 2016 compiler suite and OpenMPI 1.8.4 (with the exception of DL-POLY, which used OpenMPI 2.0.0). The NVIDIA CUDA toolkit version 7.5.18 was used for the GPU compilation and the framework was run with Python 2.7.8. The numerical experiments were carried out on the University of Bath HPC facility “Balena”. All nodes of the cluster consist of two Intel Xeon E5-2650v2 (2.6GHz) processors with eight cores each; in addition some nodes are equipped with Nvidia Tesla K20X GPU accelerator cards. As the GPU port of LAMMPS offloads the force calculation, we allowed LAMMPS to use all 16 cores of the host CPU along with the GPU. In contrast, in our framework the entire simulation is run on the GPU and it is sufficient to use a single MPI rank which acts as the host controller.

We use the parameters in Tab. 3, adapted from a LAMMPS benchmark [42]. All three codes implement the neighbour list method for force calculations. For LAMMPS and our framework the extended cutoff \bar{r}_c in Eq. (3) was chosen such that $\delta = \bar{r}_c - r_c = 0.1r_c$ with a neighbour list update every 20 iterations. In contrast, DL-POLY automatically updates the neighbour-list when necessary. The total integration time on up to 1024 cores (64 nodes) and up to 8 GPUs is plotted in Fig. 6 (left); grey regions indicate core counts contained

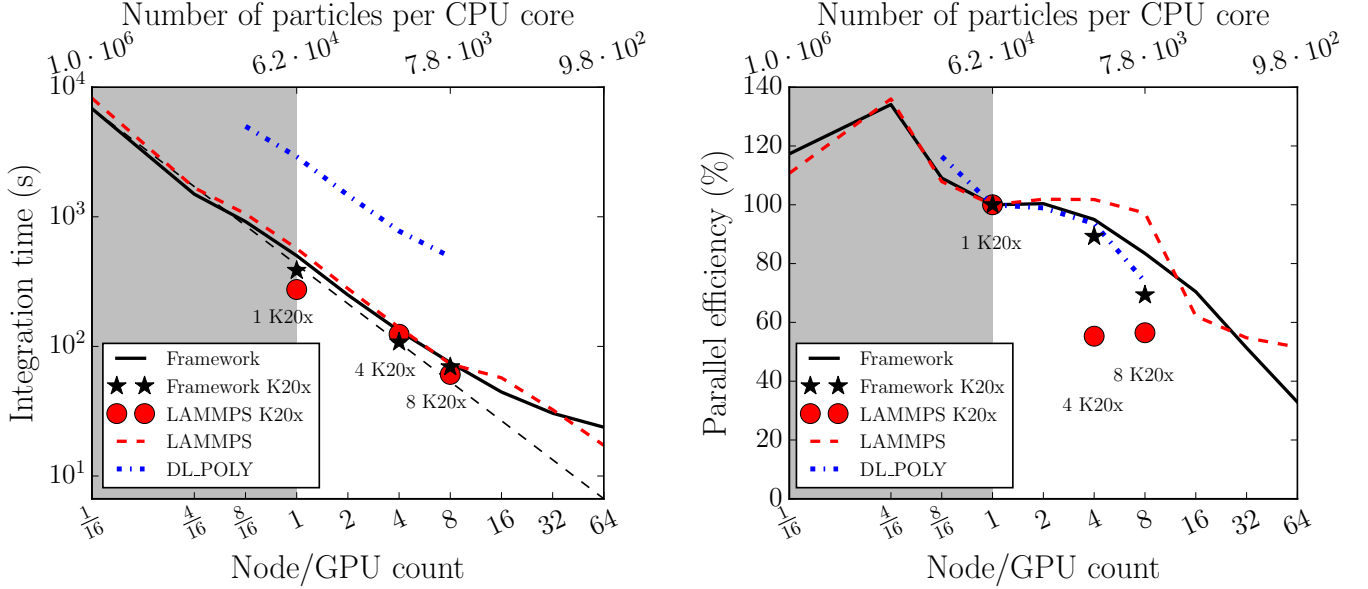


Figure 6: Strong scaling experiment: time taken to propagate $N = 10^6$ particles over $N_{\max} = 10^4$ time steps (left) and parallel efficiency (right). Efficiency relative to one full node (16 cores) is calculated according to Eqn. (10). In the left plot perfect scaling is indicated by the dashed gray line.

within a single CPU node. On the largest core count (1024 cores) the average local problem size is reduced to 1,000 particles per processor. To provide a fair comparison, one K20X GPU is compared to a full 16 cores CPU node. We write $t(p, n)$ for the measured wallclock time required to integrate a system with n particles on p CPU nodes or GPUs. The corresponding parallel efficiency (relative to one CPU node or one GPU) is defined as

$$\text{Strong parallel efficiency} = \frac{t(1, n)}{p \times t(p, n)} \quad (10)$$

and shown in Fig. 6 (right). The absolute times demonstrate that the framework provides comparable performance and scalability to DL-POLY and LAMMPS. In fact we find that for this particular setup both LAMMPS and our code are significantly faster than DL-POLY and scale better. It should be kept in mind, however, that currently both LAMMPS and DL-POLY have a much wider range of applications and provide functionality which is not yet implemented in our framework. A socket-to-socket comparison demonstrates that one full GPU can only deliver a slightly higher performance than a full CPU node. Again, the same is observed for LAMMPS. The framework can make effective use of multi-GPU systems to accelerate computation.

To test performance for very large problem sizes we also carried out a weak scaling experiment. In this setup the average work per unit computational resource is fixed and the total problem size grows proportional to the number of nodes. A system with 512,000 particles per CPU core (8,192,000 particles per node) was integrated over 5000 timesteps. For the largest computational configuration (1024 cores) the total problem size is about half a billion ($5.24 \cdot 10^8$) particles. All other system parameters are unchanged from Tab. 3. The total time for increasing problem sizes is shown in Fig. 7 (left). The weak parallel

efficiency is defined as

$$\text{Weak parallel efficiency} = \frac{t(1, n)}{t(p, n \cdot p)} \quad (11)$$

and plotted in Fig. 7 (right). We observe that (relative to one node) the parallel efficiency never drops below 90% and conclude that the framework will effectively scale to systems containing very large numbers of particles on a significant core count.

The number of particles on a single CPU node in the previous weak scaling run is too large to fit into GPU memory. To also compare the weak scalability of the generated CPU and GPU code we therefore repeat the same experiment with a reduced number of 512,000 particles per node. The resulting time and parallel efficiency are shown in Fig. 8. While the parallel efficiency is worse for the GPU, it never drops below 60%. On one node the GPU code is about twice as fast as the CPU code and on 16 nodes this speedup factor drops to around $1.3\times$. This can be explained by the fact that on one node the CPU implementation is slower and therefore communication overheads will have a relatively larger impact on the GPU code. To improve scalability further, we will investigate overlapping communication and communication in the future. This, however, is usually more challenging on GPUs due to the reduced work in halo regions.

5.1.1 Absolute performance

To quantify the absolute performance on both CPU and GPU we use data collected in the second weak scaling experiment (see Fig. 8). The computationally most expensive operation in the simulation is the force update step performed with a particle pair loop. This accounts for 54.8% of the total runtime on the CPU and 36.9% on the GPU. As in this simulation the potential energy




Figure 7: CPU-only weak scaling experiment: time taken to integrate the system over $N_{\max} = 5000$ time steps (left) and parallel efficiency (right). The efficiency relative to one full node (right) is calculated according to Eqn. (11). The top horizontal axes shows the total number N of particles in the system; the number of particles per core is kept fixed at 512,000 (8,192,000 particles per node). Perfect scaling is indicated by the horizontal dashed line.




Figure 8: CPU-GPU weak scaling experiment with reduced particle number: time taken to simulate $N_{\max} = 5000$ time steps (left) parallel efficiency relative to a single GPU/node, calculated according to Eqn. (11) (right). The number of particles per node is kept fixed at 512,000. Perfect scaling is indicated by the horizontal dashed line.

kernel	Intel Xeon node		K20X GPU	
	peak	time	peak	time
Force	16.5%	54.8%	11.9%	36.9%
Force & PE	7.5%	6.5%	14.3%	2.6%

Table 4: Absolute performance metrics (as percentage of peak performance and integration time) for two kernels recorded from GPU weak scaling experiment presented in Fig. 8.

was updated every 10 iterations, we also report performance metrics for the combined force- and potential-energy (PE) update.

With the vector instruction set each core of an E5-2650v2 (2.6 GHz) Intel CPU can perform 4 double precision additions and 4 double precision multiplications per clock cycle, resulting in a total performance of 332.8 GFLOPs per node. The peak double precision floating point performance of the nVidia Tesla K20x GPU is quoted as 1.31 TFLOPs [43].

Absolute performance numbers for a single-node run are reported in Tab. 4. The measured times only include the time spent in the auto-generated C-code, but we found that the launch of a shared library function from Python has a negligible overhead ($\approx 10\text{--}20\mu\text{s}$). Since the system is spatially homogeneous and there is little load imbalance, we report measurements collected by a single core on the fully populated node. The results demonstrate that the computationally most relevant kernels use a significant fraction of the peak floating point performance. As confirmed by the report generated by the compiler, the kernel for the Lennard-Jones force calculation in Listing 9 is automatically vectorised.

5.2 Structure analysis algorithms

We finally demonstrate how the structure analysis algorithms described in Section 4 can be implemented with our framework. For this we first add an on-the-fly implementation of the BOA analysis method. This is achieved by extending the main timestepping loop in Algorithm 6 by calls to the `PairLoop` and `ParticleLoop` which evaluate Q_ℓ according to Algorithms 1 and 2. The source code is available in the `examples/on-the-fly-analysis` subdirectory of [37].

To initialise the simulation, 125000 identical particles are arranged in a periodic cubic lattice and their velocities are sampled from a normal distribution. After allowing the system to equilibrate for 50,000 steps in an microcanonical ensemble we coupled the system to an Andersen thermostat with a target temperature near zero for 500,000 iterations. The final configuration consists of two distinct regions. The first is void of particles while the second contains a crystal structure. Fig. 9 shows the change of Q_4 , Q_5 and Q_6 throughout the simulation. A distribution of the Q_4 and Q_6 values at the final timestep is shown in Figs. 11 and 12 in B. This distribution describes the proportion of FCC and HCP in the final configuration as classified by the BOA method. In




Figure 9: Evolution of mean Q_4 , Q_5 and Q_6 values over the course of the simulation. The horizontal dashed lines plot the expected Q_4 and Q_6 values of a perfect FCC lattice.

parameter	value
Number of atoms per node:	524288
Number of time steps: N_{\max}	5000
Non-dimensionalised density: ρ	0.8442
Force cutoff: r_c	3.0
Force extended cutoff: $\bar{r}_c = r_c + \delta$	3.3
Steps between neighbour list updates:	18

Table 5: Parameters of bond order analysis weak scaling experiment. Units are chose such that $\sigma = \epsilon = 1$.

this work we purely focus on the implementation of the method and do not attempt a physical interpretation of the results.

To demonstrate that the resulting code still scales well in parallel, we carry out a weak scaling experiment with the parameters in Tab. 5. The results are shown in Fig. 10 and confirm that adding the on-the-fly analysis and thermostat have no negative impact on scalability.

Finally the common neighbour analysis was implemented as a parallel post-processing step. C-Kernels for Algorithms 3, 4, 5 and 7 can be found in the subdirectory `examples/structure/cna`. We validated our implementations by verifying that perfect crystals are correctly classified in each of the FCC, BCC and HPC configurations. We then applied the method to the test case with 125000 particles mentioned above. For the final configuration the algorithm classified 19360 (15.5%) particles as FCC and 13052 (10.4%) particles as HCP while 92588 (74.1%) particles were left unclassified. Again a physical interpretation of this result would be beyond the scope of this article.

6 Conclusions

The key computational components of a Molecular Dynamics simulation can be expressed as loops over all particles or all particle pairs. Based on this observation, we described an abstraction for implementing those loops




Figure 10: Weak scaling experiment that combines a simulation with on-the-fly analysis. Time taken to integrate 5000 steps, black dashed line indicates perfect scaling relative to a single node (left). Parallel efficiency relative to a single node (right).

and introduced the necessary data structures and execution model. Our approach is inspired by the OP2 and PyOP2 frameworks for the solution of PDEs with grid based methods. We implemented a Python-based code generation system which allows the developer to write performance portable molecular dynamics algorithms based on a separation of concerns philosophy. By considering two popular analysis methods for the classification of crystalline structures, we showed that it is easy to apply our approach to write performant and scalable analysis code. In principle the framework also allows for biasing dynamics within a simulation dependent on the local environment of each particle.

The performance and scalability of our code generation framework compares favourably to two existing and well established Molecular Dynamics codes (LAMMPS and DL-POLY) both on CPUs and GPUs. This demonstrates that for the model system considered here the code generation approach does not introduce any computational overheads; the autogenerated code runs at similar speed as monolithic codes in C++ (LAMMPS) or Fortran (DL-POLY). We stress, however, that our main aim is not to out-perform existing codes but rather explore new ways of implementing both timestepping methods and analysis algorithms with minimal programmer effort.

There are many ways in which our framework has to be extended to provide similar functionality to existing MD packages. Computationally the most serious limitation is the lack of long range force calculations. To address this issue we will introduce support for electrostatic forces either via a Ewald summation [44] or a Fast-Multipole method [45]. Since the functional form of the electrostatic interaction is fixed, this could also be achieved by incorporating a standalone C-code coupled to the code

generation framework. Other important extensions include support for constraints and multiple species. While currently different species can be simulated by adding a species label as a `ParticleDat` and adding corresponding if-branches to the computational kernels, this is clearly not efficient and should be replaced by native support in the fundamental data structures.

The performance of the GPU implementation of an algorithm is sensitive to the memory access pattern. At the beginning of a simulation particles are arranged in an ordered fashion in memory that corresponds to the physical location of the particle. As the simulation evolves the movement of particles within the simulation domain introduces an essentially random ordering of particles in memory. The results we present exhibit a slow down effect as the simulation evolves due to this sub-optimal memory ordering effect. Future versions of the framework will periodically reorder the particle data to mitigate this effect.

Finally, automatic generation of kernels from the analytical form of the potential as implemented in the OpenMM library [36] could be added. We stress, however, that it is important to still allow the user to also implement arbitrary kernels by hand to cover more general applications.

7 Acknowledgements

We would like to thank Alexander Stukowski (Darmstadt) for useful correspondence to clarify the exact definitions of the quantities in the common neighbour analysis algorithm. The PhD project of William Saunders is funded by an EPSRC studentship. This research made use of the Balena High Performance Computing (HPC)

A Kernels

This appendix lists some C-kernels which are used for the Lennard-Jones force calculation in Section 5 and in the common neighbour analysis discussed in Section 4.2. The full source code can be found in [37].

A.1 Force calculation

The Lennard-Jones potential in Eqn. (9) gives rise to the force

$$\begin{aligned} \mathbf{F}(\mathbf{r}) &= -\nabla V(r) = -\frac{\mathbf{r}}{r} \frac{\partial V}{\partial r} \\ &= \frac{48\epsilon}{\sigma^2} \mathbf{r} \left(\left(\frac{\sigma}{r} \right)^{14} - \frac{1}{2} \left(\frac{\sigma}{r} \right)^8 \right). \end{aligned} \quad (12)$$

The corresponding kernel for the force- and potential calculation is shown in Listing 9. The particle position is passed in as the `ParticleDat` `r` and the resulting force and potential energy are returned in the `ParticleDat` `F` and `ScalarArray` `u`. The squared cutoff distance r_c^2 and the numerical constants σ^2 , $C_V = 4\epsilon$ and $C_F = -48\epsilon/\sigma^2$ are passed to the pairloop as `Constant` objects. Since we use a hard cutoff, the force and potential are nonzero only if $(\mathbf{r}^{(i)} - \mathbf{r}^{(j)})^2 \leq r_c^2$ and only need to be calculated in this case. However, to ensure that the code can be vectorised, the force and potential is calculated for all relative distances $\mathbf{r}^{(i)} - \mathbf{r}^{(j)}$ and written to the variable `F` with a ternary operator.

Listing 9: Lennard-Jones kernel

```
const double dr0 = r.i[0] - r.j[0];
const double dr1 = r.i[1] - r.j[1];
const double dr2 = r.i[2] - r.j[2];
// Calculate squared distance
// dr2 = |r_i - r_j|^2
double dr_sq = dr0*dr0+dr1*dr1+dr2*dr2;
// (sigma/dr)^2
const double r_m2 = sigma2/dr_sq;
// (sigma/dr)^4
const double r_m4 = r_m2*r_m2;
// (sigma/dr)^6
const double r_m6 = r_m4*r_m2;
// (sigma/dr)^8
const double r_m8 = r_m4*r_m4;
// Increment potential energy
u[0] += (dr_sq < rc_sq) ?
    CV*((r_m6-1.0)*r_m6+0.25) : 0.0;
const double f_tmp = CF*(r_m6-0.5)*r_m8;
// Increment forces
F.i[0] += (dr_sq < rc_sq) ? f_tmp*dr0 : 0.0;
F.i[1] += (dr_sq < rc_sq) ? f_tmp*dr1 : 0.0;
F.i[2] += (dr_sq < rc_sq) ? f_tmp*dr2 : 0.0;
```

A.2 Common Neighbour analysis

Computational kernels for Algorithms 3 and 4 in the common neighbour analysis method are shown in Listings 10 and 11. The `ParticleDats` used in those kernels are related to the variables introduced in Section 4.2 and summarised in Tab. 6.

	description	ParticleDat
$\mathbf{r}^{(i)}$	particle position	<code>r</code>
$G^{(i)}$	global id	<code>id</code>
$E^{(i)}$	array repr. of $\mathcal{E}_d^{(i)} \cup \bar{\mathcal{E}}^{(i)}$	<code>bond</code>
$\nu_{nb}^{(i)}$	# of bonded neighbours	<code>n_nb</code>
$\nu_b^{(i)}$	# of bonds	<code>n_bond</code>

Table 6: Variables and `ParticleDats` used in the common neighbour analysis kernels in Listings 10 and 11.

Listing 10: CNA kernel for direct bond calculation in Algorithm 3.

```
// Calculate squared distance
const double dr0 = r.i[0] - r.j[0];
const double dr1 = r.i[1] - r.j[1];
const double dr2 = r.i[2] - r.j[2];
double dr_sq = dr0*dr0+dr1*dr1+dr2*dr2;
if (dr_sq < rc_sq) {
    // Add direct bond
    bond.i[2*n_bond.i[0]] = id.i[0];
    bond.i[2*n_bond.i[0]+1] = id.j[0];
    // Increment number of neighbours
    n_nb.i[0]++;
    // Increment number of bonds
    n_bond.i[0]++;
}
```

Listing 11: CNA kernel for direct bond calculation in Algorithm 4.

```
// Calculate squared distance
const double dr0 = r.i[0] - r.j[0];
const double dr1 = r.i[1] - r.j[1];
const double dr2 = r.i[2] - r.j[2];
double dr_sq = dr0*dr0+dr1*dr1+dr2*dr2;
if (dr_sq < rc_sq) {
    for (int k=0; k<n_nb.j[0]; ++k) {
        // Add indirect bond
        if (bond.j[2*k+1] != id.i[0]) {
            bond.i[2*n_bond.i[0]] = bond.j[2*k];
            bond.i[2*n_bond.i[0]+1] =
                bond.j[2*k+1];
            // Increment number of bonds
            n_bond.i[0]++;
        }
    }
}
```

B Bond Order Analysis

Figures 11 and 12 show the final distribution of the order parameters Q_4 and Q_6 in the numerical experiment described in Section 5.2.

C Largest subcluster algorithm

Algorithm 7 can be used to calculate the size of the largest connected component of a graph given by a set of edges \mathcal{E} . For this the edges in each subgraph are counted with a breadth-first traversal, counting and removing all visited edges in the process.




Figure 11: Probability density of Q_4 values in final system configuration. Dashed vertical line at $Q_4 = 0.097$ is the expected Q_4 value of a perfect hcp lattice. Dashed vertical line at $Q_4 = 0.191$ is the expected Q_4 value of a perfect fcc lattice.




Figure 12: Probability density of Q_6 values in final system configuration. Dashed vertical line at $Q_6 = 0.485$ is the expected Q_6 value of a perfect hcp lattice. Dashed vertical line at $Q_6 = 0.575$ is the expected Q_6 value of a perfect fcc lattice.

Algorithm 7 Calculate maximal cluster size.

Input: graph defined by a set of edges \mathcal{E} .

Output: S_{\max} , the size of the largest cluster

```

1:  $S_{\max} \mapsto 0$ 
2: while  $\mathcal{E} \neq \emptyset$  do
3:    $S \mapsto 0$ 
4:   Pick some edge  $(v_1, v_2) \in \mathcal{E}$ 
5:    $\mathcal{Q} \mapsto \{v_1\}$ 
6:   while  $\mathcal{Q} \neq \emptyset$  do
7:     Pick some  $v \in \mathcal{Q}$  and remove it from  $\mathcal{Q}$ 
8:      $\mathcal{P} \mapsto \{(v, w) \in \mathcal{E}\}$ 
9:      $\mathcal{Q} \mapsto \mathcal{Q} \cup \{w : (v, w) \in \mathcal{P}\}$ 
10:     $S \mapsto S + |\mathcal{P}|$ 
11:    Remove all edges  $e \in \mathcal{P}$  from  $\mathcal{E}$ 
12:   end while
13:    $S_{\max} \mapsto \max\{S, S_{\max}\}$ 
14: end while

```

References

- [1] M. T. Nelson, W. Humphrey, A. Gursoy, A. Dalke, L. V. Kalé, R. D. Skeel, K. Schulten, **NAMD: a parallel, object-oriented molecular dynamics program**, International Journal of High Performance Computing Applications 10 (4) (1996) 251–268. doi:10.1177/109434209601000401. URL <https://doi.org/10.1177/109434209601000401>
- [2] J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, K. Schulten, **Scalable molecular dynamics with NAMD**, Journal of Computational Chemistry 26 (16) (2005) 1781–1802. doi:10.1002/jcc.20289. URL <https://doi.org/10.1002/jcc.20289>
- [3] S. Plimpton, **Fast parallel algorithms for short-range molecular dynamics**, Journal of Computational Physics 117 (1) (1995) 1 – 19. doi:10.1006/jcph.1995.1039. URL <https://doi.org/10.1006/jcph.1995.1039>
- [4] H. Berendsen, D. van der Spoel, R. van Drunen, **GROMACS: A message-passing parallel molecular dynamics implementation**, Computer Physics Communications 91 (1) (1995) 43 – 56. doi:10.1016/0010-4655(95)00042-E. URL [https://doi.org/10.1016/0010-4655\(95\)00042-E](https://doi.org/10.1016/0010-4655(95)00042-E)
- [5] S. Pronk, S. Páll, R. Schulz, P. Larsson, P. Bjelkmar, R. Apostolov, M. R. Shirts, J. C. Smith, P. M. Kasson, D. van der Spoel, B. Hess, E. Lindahl, **GROMACS 4.5: a high-throughput and highly parallel open source molecular simulation toolkit**, Bioinformatics 29 (7) (2013) 845–854. doi:10.1093/bioinformatics/btt055. URL <https://doi.org/10.1093/bioinformatics/btt055>
- [6] W. Smith, T. Forester, **DL_POLY_2.0: A general-purpose parallel molecular dynamics simulation package**, Journal of Molecular Graphics 14 (3) (1996) 136 – 141. doi:10.1016/S0263-7855(96)00043-4. URL [http://dx.doi.org/10.1016/S0263-7855\(96\)00043-4](http://dx.doi.org/10.1016/S0263-7855(96)00043-4)
- [7] I. T. Todorov, W. Smith, K. Trachenko, M. T. Dove, **DL_POLY_3: new dimensions in molecular dynamics simulations via massive parallelism**, J. Mater. Chem. 16 (2006) 1911–1918. doi:10.1039/B517931A. URL <http://dx.doi.org/10.1039/B517931A>
- [8] M. Karplus, G. A. Petsko, **Molecular dynamics simulations in biology**, Nature 347 (6294) (1990) 631–639. doi:10.1038/347631a0. URL <https://doi.org/10.1038/347631a0>

- [9] D. Rapaport, **Self-assembly of polyhedral shells: a molecular dynamics study**, *Physical Review E* 70 (5) (2004) 051905. doi:10.1103/PhysRevE.70.051905. URL <https://doi.org/10.1103/PhysRevE.70.051905>
- [10] C. Horowitz, J. Hughto, A. Schneider, D. Berry, **Neutron star crust and molecular dynamics simulation**, arXiv preprint arXiv:1109.5095. URL <https://arxiv.org/abs/1109.5095>
- [11] N. R. Williams, M. Molinari, S. C. Parker, M. T. Storr, **Atomistic investigation of the structure and transport properties of tilt grain boundaries of UO₂**, *Journal of Nuclear Materials* 458 (2015) 45 – 55. doi:10.1016/j.jnucmat.2014.11.120. URL <https://doi.org/10.1016/j.jnucmat.2014.11.120>
- [12] N. Isaac, *The mathematical principles of natural philosophy*, Book II.
- [13] P. J. Steinhardt, D. R. Nelson, M. Ronchetti, **Bond-orientational order in liquids and glasses**, *Phys. Rev. B* 28 (1983) 784–805. doi:10.1103/PhysRevB.28.784. URL <https://doi.org/10.1103/PhysRevB.28.784>
- [14] J. D. Honeycutt, H. C. Andersen, **Molecular dynamics study of melting and freezing of small Lennard-Jones clusters**, *The Journal of Physical Chemistry* 91 (19) (1987) 4950–4963. doi:10.1021/j100303a014. URL <https://doi.org/10.1021/j100303a014>
- [15] A. Stukowski, **Structure identification methods for atomistic simulations of crystalline materials**, *Modelling and Simulation in Materials Science and Engineering* 20 (4) (2012) 045021. URL <http://stacks.iop.org/0965-0393/20/i=4/a=045021>
- [16] J. E. Stone, J. C. Phillips, P. L. Freddolino, D. J. Hardy, L. G. Trabuco, K. Schulten, **Accelerating molecular modeling applications with graphics processors**, *Journal of Computational Chemistry* 28 (16) (2007) 2618–2640. doi:10.1002/jcc.20829. URL <https://doi.org/10.1002/jcc.20829>
- [17] J. A. Anderson, C. D. Lorenz, A. Travesset, **General purpose molecular dynamics simulations fully implemented on graphics processing units**, *Journal of Computational Physics* 227 (10) (2008) 5342 – 5359. doi:10.1016/j.jcp.2008.01.047. URL <https://doi.org/10.1016/j.jcp.2008.01.047>
- [18] W. M. Brown, P. Wang, S. J. Plimpton, A. N. Tharrington, **Implementing molecular dynamics on hybrid high performance computers - short range forces**, *Computer Physics Communications* 182 (4) (2011) 898 – 911. doi:10.1016/j.cpc.2010.12.021. URL <https://doi.org/10.1016/j.cpc.2010.12.021>
- [19] W. M. Brown, A. Kohlmeyer, S. J. Plimpton, A. N. Tharrington, **Implementing molecular dynamics on hybrid high performance computers - Particle-particle particle-mesh**, *Computer Physics Communications* 183 (3) (2012) 449 – 459. doi:10.1016/j.cpc.2011.10.012. URL <https://doi.org/10.1016/j.cpc.2011.10.012>
- [20] M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, E. Lindahl, **GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers**, *SoftwareX* 1-2 (2015) 19 – 25. doi:10.1016/j.softx.2015.06.001. URL <https://doi.org/10.1016/j.softx.2015.06.001>
- [21] J. Glaser, T. D. Nguyen, J. A. Anderson, P. Lui, F. Spiga, J. A. Millan, D. C. Morse, S. C. Glotzer, **Strong scaling of general-purpose molecular dynamics simulations on GPUs**, *Computer Physics Communications* 192 (2015) 97 – 107. doi:10.1016/j.cpc.2015.02.028. URL <https://doi.org/10.1016/j.cpc.2015.02.028>
- [22] C. Bertolli, A. Betts, G. Mudalige, M. Giles, P. Kelly, **Design and Performance of the OP2 Library for Unstructured Mesh Applications**, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 191–200. doi:10.1007/978-3-642-29737-3_22. URL https://doi.org/10.1007/978-3-642-29737-3_22
- [23] M. B. Giles, G. R. Mudalige, B. Spencer, C. Bertolli, I. Reguly, **Designing OP2 for GPU Architectures**, *J. Parallel Distrib. Comput.* 73 (11) (2013) 1451–1460. doi:10.1016/j.jpdc.2012.07.008. URL <http://dx.doi.org/10.1016/j.jpdc.2012.07.008>
- [24] M. B. Giles, G. R. Mudalige, Z. Sharif, G. Markall, P. H. Kelly, **Performance Analysis of the OP2 Framework on Many-core Architectures**, *SIGMETRICS Perform. Eval. Rev.* 38 (4) (2011) 9–15. doi:10.1145/1964218.1964221. URL <https://doi.org/10.1145/1964218.1964221>
- [25] I. Z. Reguly, E. László, G. R. Mudalige, M. B. Giles, **Vectorizing unstructured mesh computations for many-core architectures**, *Concurrency and Computation: Practice and Experience* 28 (2) (2016) 557–577, cpe.3621. doi:10.1002/cpe.3621. URL <http://dx.doi.org/10.1002/cpe.3621>
- [26] T. Gysi, O. Fuhrer, C. Osuna, B. Cumming, T. Schulthess, **STELLA: A domain-specific**

- embedded language for stencil codes on structured grids, in: EGU General Assembly Conference Abstracts, Vol. 16 of EGU General Assembly Conference Abstracts, 2014, p. 8464. doi:10.1145/2807591.2807627. URL <https://doi.org/10.1145/2807591.2807627>
- [27] F. Rathgeber, G. R. Markall, L. Mitchell, N. Lorian, D. A. Ham, C. Bertolli, P. H. J. Kelly, **PyOP2: A High-Level Framework for Performance-Portable Simulations on Unstructured Meshes**, in: High Performance Computing, Networking Storage and Analysis, SC Companion:, IEEE Computer Society, Los Alamitos, CA, USA, 2012, pp. 1116–1123. doi:10.1109/SC.Companion.2012.134. URL <https://doi.org/10.1109/SC.Companion.2012.134>
- [28] D. Rapaport, **Enhanced molecular dynamics performance with a programmable graphics processor**, Computer Physics Communications 182 (4) (2011) 926 – 934. doi:10.1016/j.cpc.2010.12.029. URL <https://doi.org/10.1016/j.cpc.2010.12.029>
- [29] L. Verlet, **Computer "Experiments" on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules**, Phys. Rev. 159 (1967) 98–103. doi:10.1103/PhysRev.159.98. URL <https://doi.org/10.1103/PhysRev.159.98>
- [30] M. P. Allen, D. J. Tildesley, Computer simulation of liquids, Oxford University Press, Oxford, 1989.
- [31] D. Frenkel, B. Smit, Understanding molecular simulation: from algorithms to applications.
- [32] Z. Jia, B. Leimkuhler, **Molecular simulation in the canonical ensemble and beyond**, ESAIM: Mathematical Modelling and Numerical Analysis 41 (2) (2007) 333–350. doi:10.1051/m2an:2007019. URL <https://doi.org/10.1051/m2an:2007019>
- [33] B. Leimkuhler, E. Noorizadeh, O. Penrose, **Comparing the efficiencies of stochastic isothermal molecular dynamics methods**, Journal of Statistical Physics 143 (5) (2011) 921–942. doi:10.1007/s10955-011-0210-2. URL <https://doi.org/10.1007/s10955-011-0210-2>
- [34] M. Bonomi, D. Branduardi, G. Bussi, C. Camilloni, D. Provasi, P. Raiteri, D. Donadio, F. Marinelli, F. Pietrucci, R. A. Broglia, M. Parrinello, **Plumed: A portable plugin for free-energy calculations with molecular dynamics**, Computer Physics Communications 180 (10) (2009) 1961 – 1972. doi:10.1016/j.cpc.2009.05.011. URL <https://doi.org/10.1016/j.cpc.2009.05.011>
- [35] I. Bethune, E. Breitmoser, B. Leimkuhler, **Molecular Integration Simulation Toolkit - interfacing novel integrators with Molecular Dynamics codes**, talk at ISQBP President’s Meeting, Bergen, Norway, (MIST available at <https://bitbucket.org/extasy-project/mist/wiki/About>) (2016).
- [36] P. Eastman, M. S. Friedrichs, J. D. Chodera, R. J. Radmer, C. M. Bruns, J. P. Ku, K. A. Beauchamp, T. J. Lane, L.-P. Wang, D. Shukla, T. Tye, M. Houston, T. Stich, C. Klein, M. R. Shirts, V. S. Pande, **OpenMM 4: A Reusable, Extensible, Hardware Independent Library for High Performance Molecular Simulation**, Journal of Chemical Theory and Computation 9 (1) (2013) 461–469, pMID: 23316124. doi:10.1021/ct300857j. URL <https://doi.org/10.1021/ct300857j>
- [37] W. R. Saunders, Framework snapshot, <https://doi.org/10.5281/zenodo.496142> (2017).
- [38] D. Rapaport, **Large-scale molecular dynamics simulation using vector and parallel computers**, Computer Physics Reports 9 (1) (1988) 1 – 53. doi:10.1016/0167-7977(88)90014-7. URL [https://doi.org/10.1016/0167-7977\(88\)90014-7](https://doi.org/10.1016/0167-7977(88)90014-7)
- [39] D. Rapaport, The Art of Molecular Dynamics Simulation, Cambridge University Press, Cambridge, 2004.
- [40] W. Mickel, S. C. Kapfer, G. E. Schröder-Turk, K. Mecke, **Shortcomings of the bond orientational order parameters for the analysis of disordered particulate matter**, The Journal of Chemical Physics 138 (4) (2013) 044501. doi:10.1063/1.4774084. URL <https://doi.org/10.1063/1.4774084>
- [41] W. R. Saunders, Paper data repository, <https://doi.org/10.5281/zenodo.496147> (2017).
- [42] S. C. et al, **Lennard-Jones liquid benchmark**, <http://lammps.sandia.gov/bench.html#lj>, [Online; accessed 03/06/2016] (2016).
- [43] NVIDIA Corporation, NVIDIA Tesla GPU Accelerators, <http://www.nvidia.com/content/tesla/pdf/nvidia-tesla-kepler-family-datasheet.pdf> (2017).
- [44] P. P. Ewald, **Die Berechnung optischer und elektrostatischer Gitterpotentiale**, Annalen der Physik 369 (3) (1921) 253–287. doi:10.1002/andp.19213690304. URL <https://doi.org/10.1002/andp.19213690304>
- [45] L. Greengard, V. Rokhlin, **A Fast Algorithm for Particle Simulations**, J. Comput. Phys. 73 (2) (1987) 325–348. doi:10.1016/0021-9991(87)90140-9. URL [http://dx.doi.org/10.1016/0021-9991\(87\)90140-9](http://dx.doi.org/10.1016/0021-9991(87)90140-9)